

Knowledge-Based Temporal Abstraction for Diabetic Monitoring

Yuval Shahar, Amar K. Das, Samson W. Tu, Frederic B. Kraemer, Mark A. Musen

Section on Medical Informatics, Stanford University School of Medicine
Stanford, CA 94305-5479

Abstract

We have developed a general method that solves the task of creating abstract, interval-based concepts from time-stamped clinical data. We refer to this method as knowledge-based temporal-abstraction (KBTA). In this paper, we focus on the knowledge representation, acquisition, maintenance, reuse and sharing aspects of the KBTA method. We describe five problem-solving mechanisms that solve the five subtasks into which the KBTA method decomposes its task, and four types of knowledge necessary for instantiating these mechanisms in a particular domain. We present an example of instantiating the KBTA method in the clinical area of monitoring insulin-dependent-diabetes patients.

1. THE TEMPORAL-ABSTRACTION TASK

Clinical data, such as blood-glucose values, are typically gathered over time within the scope of one or more interpretation contexts (e.g., a healthy person, an insulin-dependent diabetes patient, pre-breakfast periods, regular insulin therapy). The **temporal-abstraction (TA) task** accepts as input time-stamped parameters (e.g., blood glucose values) and events (e.g., NPH insulin injections), and returns as output interval-based, context-specific parameters at the same or at a higher level of abstraction. The process of summarizing large amounts of clinical data over time supports a physician assessing a patient's condition by creating abstract concepts (e.g., 2 weeks of LOW pre-breakfast blood glucose and HIGH pre-supper values) from raw numerical data (e.g., pre- and post-prandial blood glucose values).

The goal of the TA task is to evaluate and summarize the state of the patient over a time interval, to identify various possible problems, to assist in a revision of an existing therapy plan, or to support a generation of a new plan. In addition, generating clinically meaningful interval-based concepts supports the task of explaining a decision-support system's plans and actions to different

users (e.g., a resident physician, a nurse, an experienced clinical expert). Finally, clinical guidelines can represent goals and policies as temporal patterns to be achieved or avoided.

Several issues need to be handled by a method solving the TA task: (1) the arriving input or the queried output parameter values might be of different types (e.g., numbers, symbols) and abstraction levels (e.g., BLOOD GLUCOSE LEVEL = 64 mg%; GLUCOSE STATE = LOW); (2) input data might arrive out of temporal order, and existing interpretations should be revised accordingly; (3) several alternate interpretations might need to be maintained and followed over time; (4) from the knowledge-representation aspect, **acquisition** of necessary knowledge from domain experts should be facilitated, as well as **maintenance** of that knowledge. **Reusing** the domain-independent abstraction knowledge for solving the TA task in other domains should be possible, as well as **sharing** some of the domain-specific knowledge with other tasks in the same domain.

2. THE KNOWLEDGE-BASED TEMPORAL-ABSTRACTION METHOD

Generalizing our previous work [1,2], we have defined a domain-independent problem-solving method [3] for interpreting data in the time-oriented, knowledge-intensive domains common to clinical applications. We propose a highly modular approach, with semantics clearly defined for both the problem-solving method and the domain-specific knowledge needed by it. The **knowledge-based temporal-abstraction (KBTA)** method decomposes the TA task into five parallel subtasks (Figure 1): (1) **temporal-context restriction**: creation of relevant interpretation contexts crucial for focusing and limiting the scope of the inference, (2) **vertical temporal inference**: inference from contemporaneous propositions into higher-level concepts, (3) **horizontal temporal inference**: inference from propositions of similar type, attached to intervals that cover different time periods,

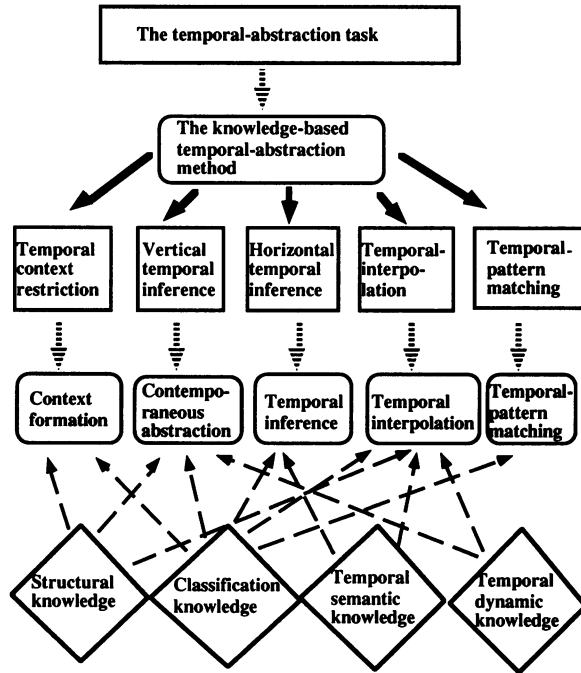


Figure 1: The knowledge-based temporal-abstraction method and its mechanisms. \square = task; oval = method or mechanism; \diamond = knowledge type; \rightarrow = USED-BY relation; \dashrightarrow = DECOMPOSED-INTO relation; $\cdots\rightarrow$ = SOLVED-BY relation.

(4) **temporal interpolation**: bridging gaps between disjoint points or intervals, associated with propositions of similar type, to create longer intervals, and (5) **temporal-pattern matching** (creation of intervals by matching of patterns over disjoint intervals, associated with propositions of various-types). An analysis of most existing temporal-reasoning systems in clinical medicine supports this decomposition [3].

The five subtasks that the KBTA *method* poses are solved, respectively, by five lower-level temporal-abstraction **mechanisms** (computational modules that cannot be decomposed further). The mechanisms include three basic TA mechanisms which we described previously [1, 2], a mechanism for matching temporal patterns, and a mechanism for creating relevant temporal interpretation contexts. Types of output abstractions include **state** (e.g., LOW) **gradient** (e.g., INCREASING) **rate** (e.g., FAST) or **pattern** (e.g., CRESCENDO). The **contemporaneous abstraction** mechanism abstracts one or more parameters and their values, attached to contemporaneous time points or time intervals, into a value of a new, abstract parameter. The **temporal inference** mechanism infers specific types of interval-based logical conclusions, given interval-

based propositions, using an extension of Shoham's temporal semantic properties [4]. Temporal inference also determines the domain value of an abstraction created from two joined abstractions (e.g., DECREASING and SAME might be concatenated into NONINCREASING). The **temporal interpolation** mechanism bridges gaps between temporal points or intervals, using domain-specific dynamic-change knowledge about the parameters involved. In particular, it uses local (forward and backward, around a time point or interval) and global (between 2 time points or intervals) **truth-persistence** functions to model a belief in the value of a $\langle \text{parameter}, \text{value}, \text{context} \rangle$ proposition [3]. Global persistence functions are represented as maximal-temporal-gap thresholds that can be bridged. The **temporal-pattern-matching** mechanism matches predefined or online queries for temporal patterns that are defined in terms of the interval-based abstractions, at any level of abstraction, created by the other TA mechanisms. The output is a higher-level parameter of the pattern abstraction type, such as REBOUND HYPERGLYCEMIA. The **context-forming** mechanism creates temporal **interpretation-context intervals** that are a temporal frame of reference for interpretation, and thus enable a TA mechanism to conclude abstractions relevant to that and only that context. The relation between an interpretation context or subcontext and its generating task, event, abstraction or supercontext can be any of Allen's 13 temporal-interval relations [5]. Thus, contexts generated by events and by abstractions also enable anticipation of future complications and interpretation of past findings in the light of the present interpretation. Creating contexts requires knowledge about the structure of clinical tasks, events, and abstractions.

3. DOMAIN-SPECIFIC KNOWLEDGE: ONTOLOGIES

To be useful for a particular clinical domain, the TA mechanisms require instantiation with domain-specific knowledge. This domain-specific knowledge, mostly declarative, is the *only* interface between the KBTA method and the knowledge engineer or the domain expert. Thus, the development of a TA system particular to a new domain relies only on creating or editing a predefined set of knowledge categories. As shown in Figure 1, we distinguish among four domain **knowledge types** used by the TA mechanisms: (1) **structural knowledge** (e.g., IS-A and PART-OF relations in the domain); (2) **classification knowledge** (e.g., classification of blood glucose value ranges into HYPOGLYCEMIA, LOW, NORMAL, HIGH); (3) **temporal semantic knowledge** (e.g., the

relations among propositions attached to intervals and their subintervals); and (4) **temporal dynamic knowledge** (e.g., persistence of the value of a parameter over time).

The domain-specific knowledge required by the TA mechanisms is represented as a **parameter-properties ontology**—a theory that represents the raw and abstract parameters in that domain (e.g., blood glucose value and state abstractions), their temporal properties, and the relations among them (e.g., IS-A, ABSTRACTED-INTO) [2]. The parameter-properties ontology is used by all the TA mechanisms. The context-forming mechanism refers also to an ontology of **events** such as insulin administration, and an ontology of **interpretation contexts**.

4. THE RÉSUMÉ SYSTEM AND THE DIABETES DOMAIN

We have developed a software system, **RÉSUMÉ**, that implements the temporal-abstraction method [2]. A simple TA pattern-matching language queries the internal temporal fact base for particular predefined temporal patterns or for online interaction with the user. More complex queries can be answered by a relational database temporal-query system, **Chronus**, that is an extension of the temporal pattern-matching mechanism [6]. The TA mechanisms do not operate in a fixed order; they are activated by the currently available data and the previously derived abstractions. In addition, an underlying truth-maintenance system updates the temporal-interval conclusions, since these are by nature **nonmonotonic** and therefore **defeasible**, that is, their validity depends on primitive data that might be modified when more past or present data are known. The control structure implemented in the RÉSUMÉ system is specialized to the TA task, and allows several levels of task-specific control (e.g., desired output types and abstraction classes, which TA mechanisms to use, the relevant temporal contexts). It is thus both a data- and a goal-driven task-specific control.

We applied the RÉSUMÉ methodology to the area of treating insulin-dependent **diabetes mellitus (DM)** patients. One of us (F.B.K.) is a diabetologist, and was the domain expert for this experiment. We created a parameter-properties ontology (Figure 2), an event ontology (Figure 3), and a context ontology (Figure 4). Acquiring the three initial core ontologies required two meetings of 2 hours each. Administrations of regular insulin and of isophane insulin suspension (NPH) are *events*, generating different insulin-action *interpretation*

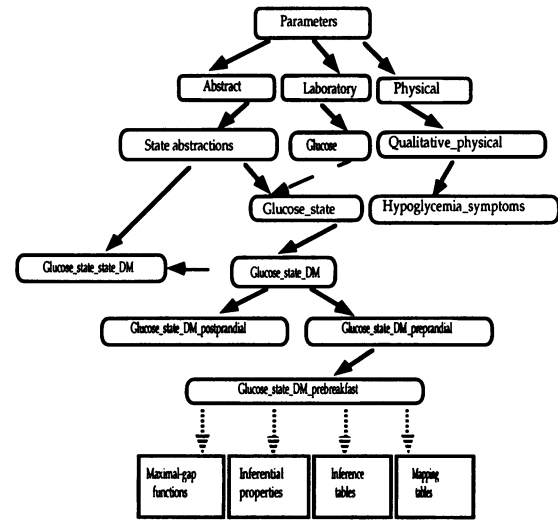


Figure 2: Part of the diabetes parameter-properties ontology. ○ = class; □ = property; → = IS-A relation; = PROPERTY-OF relation; —→ = ABSTRACTED-INTO relation.

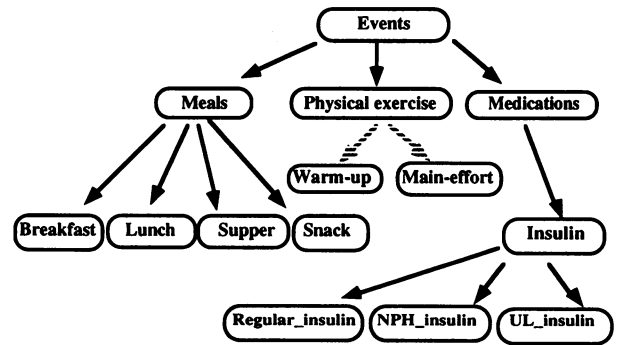


Figure 3: Part of the diabetes event ontology. ○ = class; → = IS-A relation; = PART-OF relation.

contexts that are *subcontexts* of the DM *interpretation context*. Meal events create pre- and post-prandial contexts—Glucose_state_DM_prebreakfast values [See Figure 2] can thus be inferred regardless of absolute time. The Glucose_state abstract parameter has six values that correspond to the ranges used by the domain expert (HYPOGLYCEMIA, LOW, NORMAL, HIGH, VERY HIGH, EXTREMELY HIGH). These values are sensitive to the context in which they are generated; for instance, postprandial values allow for a higher range of the normal value. Glucose_state value propositions in the same DM context have the semantic property of being **concatenable** into propositions holding over longer intervals [4]; *same-day* values between different

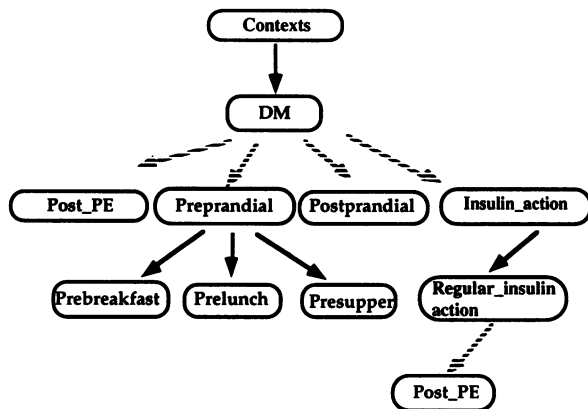


Figure 4. Part of the diabetes context ontology. = class; = IS-A relation; = SUBCONTEXT relation.

preprandial phases can be bridged up to 6 to 8 hours apart (defined by a truth-persistence interpolation function). *Same-phase* parameters such as *Glucose_state_DM_prebreakfast* have longer global persistence, since they are typically bridged over 24 to 28 hours, using another interpolation function. A higher-level abstraction of the state of glucose, *Glucose_state_state*, maps its six values into three categories (LOW, NORMAL, HIGH, or L, N, H, for short), has different semantic properties, and allows creation of daily pattern abstractions such as *LLH* (e.g., from prebreakfast, prelunch and presupper glucose values, respectively). Recognition of such patterns can be highly useful when deciding how to modify a patient's insulin regimen; noting their prevalence is an important step in determining if the pattern is a common one for the patient. Asserting anywhere in the temporal fact base an event named *DM_planning* initiates the reasoning by generating a retrospective DM interpretation context for the preceding 2 weeks (this time window is used by the domain expert in practice and is modifiable) that enables creation of the DM domain abstractions.

We applied the *RÉSUMÉ* system to electronic data from insulin-dependent-diabetes patients. The input to *RÉSUMÉ* included both the diabetes ontology (figures 2 through 4) and the patient-specific raw data. A sample of the results is shown in Figure 5. In this particular time window, two significant findings are highlighted: The *Glucose_state_state* parameter in the presupper context had the value *HIGH* for a period of more than 3 days, and a diurnal pattern of *NORMAL* or *LOW* blood glucose levels at morning and lunch, and *HIGH* pre-supper glucose levels (e.g., *NNH*, *NLH*)

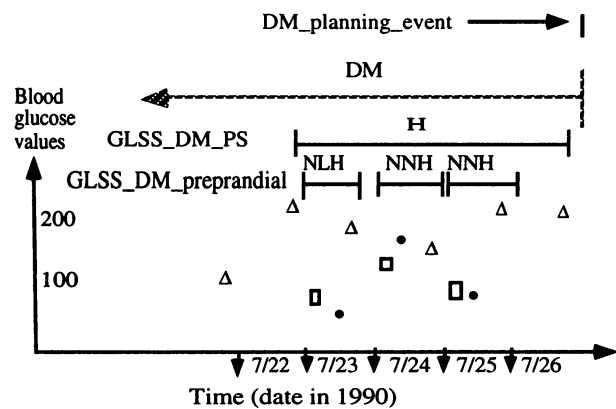


Figure 5: An example of abstraction by the *RÉSUMÉ* system of data from patient 3. = (open) context interval; = abstracted interval; = prebreakfast glucose; = prelunch glucose; = presupper glucose; *GLSS_DM_PS* = *Glucose_state_state* abstraction in the DM and pre-supper context; *GLSS_DM_PREPRANDIAL* = *Glucose_state_state* abstraction in the DM and preprandial context.

appeared at least three times in the same week. The combined pattern suggests an adjustment of the intermediate-acting insulin (e.g., *NPH*). The pattern can be predefined as an internal pattern-type parameter or can be noted in response to an external online query for state abstractions of the *Glucose_state* parameter.

5. DISCUSSION

The truth-maintenance system in *RÉSUMÉ* resembles Russ's **temporal control structure (TCS)** system [7]. However, TCS leaves all domain-specific temporal reasoning to the user-created procedures. In contrast, the *RÉSUMÉ* domain-independent (but specific to the TA task) TA mechanisms perform all of the TA, given the declarative representation of the domain's ontology.

The **TrenDx** system of Haimowitz and Kohane [8] builds on Kohane's constraint-satisfaction **temporal-utilities package** [9], and defines domain-specific patterns called **trend templates (TTs)**. *TrenDx* is useful in detecting that the data is consistent with one or more TTs, including TTs of which only a part is observed. The goal of *TrenDx* is different from that of *RÉSUMÉ*. *TrenDx* does not create any intermediate abstractions, since its goal is not to abstract, summarize, or answer queries about the data, as it is in the TA task, but rather to match data efficiently against a set of predefined patterns. Data can only be accepted

at the lowest level; thus, no input of intermediate-level abstractions is possible. No explicit domain ontology of parameters and events exists, and a constraint (e.g., significant change in a parameter) might be repeated with the same implicit role in different TTs and even at different parts of the same TT. Like RÉSUMÉ, Trendx assumes implicitly an ill-defined domain that cannot be modeled easily numerically, and therefore requires detection of essentially associative temporal patterns.

Kahn's TOPAZ system [10] integrates a quantitative physiological model and a symbolic model for aggregation of clinically significant intervals. TOPAZ can associate interpretation methods with an interval representing a *context* of interest. RÉSUMÉ extends this capability by the context-forming mechanism, which uses an explicit context ontology to enable creation of context-specific abstractions and activation of specific functions, but does not limit generated interpretation contexts to the temporal extent of the parent event, allowing any desired relation between the generating interval and the generated context. Lehmann's AIDA system [11] is a diabetes-treatment decision support prototype system, whose underlying model attempts to reflect the (patho)physiology of insulin action and carbohydrate absorption in quantitative terms. Note that systems such as TOPAZ and AIDA assume a precise underlying mathematical model of the domain; most clinical domains defy complete quantitative modeling.

It might be desirable to detect patterns defined by *events*, such as insulin use, and not only by *parameters*, such as glucose states. Such patterns might generate more meaningful interpretation contexts. Such work has been described by Kahn and his colleagues [12] with encouraging results for an algorithm combining clinical and temporal considerations.

Acknowledgments

Dr. M. G. Kahn of the University of Washington, St. Louis, supplied the data. This work has been supported in part by grant HS06330 from the Agency for Health Care Policy and Research, by grants LM05157 and LM05305 from the National Library of Medicine, and by a gift from Digital Equipment Corporation. Dr. Musen is a recipient of National Science Foundation Young Investigator Award IRI-9257578.

References

[1]. Shahar, Y., Tu, S.W., and Musen, M.A. (1992). Knowledge acquisition for temporal-abstraction mechanisms. *Knowledge Acquisition* 4, 217–236.

[2]. Shahar, Y., and Musen, M.A. (1993). RÉSUMÉ: A temporal-abstraction system for patient monitoring. *Computers and Biomedical Research* 26, 255–273.

[3]. Shahar, Y. [1994]. *A Knowledge-Based Method for Temporal Abstraction of Clinical Data*. Ph.D. Dissertation, Program in Medical Information Sciences, Stanford University, CA.

[4]. Shoham, Y. (1987). Temporal logics in AI: Semantical and ontological considerations. *Artificial Intelligence* 33, 89–104.

[5]. Allen, J.F. (1984). Towards a general theory of action and time. *Artificial Intelligence* 23, 123–154.

[6]. Das, A.K., and Musen, M.A. (in press). A temporal query system for protocol-directed decision support. *Methods of Information in Medicine*.

[7]. Russ, T.A. (1989). Using hindsight in medical decision making. *Proceedings, Thirteenth Annual Symposium on Computer Applications in Medical Care* (L. C. Kingsland, Ed.), pp. 38–44, Washington, D.C: IEEE Comput. Soc. Press.

[8]. Haimowitz, I.J., and Kohane, I.S. (1993). An epistemology for clinically significant trends. *Proceedings of the Tenth National Conference on Artificial Intelligence*, pp. 176–181, Menlo Park, CA: AAAI Press.

[9]. Kohane, I.S. (1987). Temporal reasoning in medical expert systems. *Technical Report 389, Laboratory of Computer Science, Massachusetts Institute of technology*, Cambridge, MA.

[10]. Kahn, M.G. (1991). Combining physiologic models and symbolic methods to interpret time-varying patient data. *Methods of information in Medicine* 30, 167–178.

[11]. Lehmann, E.D., Deutsch, T., Carson, E.R., and Sonksen, P.H. (1994). AIDA: an interactive diabetes advisor. *Computer Methods and Programs in Biomedicine* 41, 183–203.

[12]. Kahn, M.G., Abrams, C.A., Cousins, S.B., Beard, J.C., and Frisse, M.E. (1990). Automated interpretation of diabetes patient data: Detecting temporal changes in insulin therapy. *Proceedings, Fourteenth Annual Symposium on Computer Applications in Medical Care* (R. A. Miller, Ed.), pp. 569–573, IEEE Comput. Soc. Press, Los Alamitos.